

Unification of Maximum Entropy and Bayesian Inference via Plausible Reasoning

Alexis Akira Toda ^{*†}

This Version: January 14, 2013

Abstract

This paper modifies Jaynes’s axioms of plausible reasoning and derives the minimum relative entropy principle, Bayes’s rule, as well as maximum likelihood from first principles. The new axioms, which I call the *Optimum Information Principle*, is applicable whenever the decision maker is given the data and the relevant background information. These axioms provide an answer to the question “why maximize entropy when faced with incomplete information?”

1 Introduction

Bayesian inference [1] and the maximum entropy principle (MaxEnt) of Jaynes [9] are valid methods of inference when the decision maker is faced with incomplete information. Although these methodologies are quite distinct, they often give similar results. A few authors have hinted at the possibility of deriving both methods from first principles. For instance, as the sample size increases, [23] showed that the distribution of a random variable conditional on empirical moment constraints (computed by Bayes’s rule) converges to the minimum relative entropy distribution subject to the same population moment constraints. Conversely, [24] showed that Bayes’s rule can be derived from a variational principle of information processing.

One possibility of deriving *both* the maximum entropy principle and Bayes’s rule is to axiomatize plausible reasoning, as [14, 13, 3, 12] attempted. In the most primitive form, Jaynes [12] suggested desiderata that should be employed in plausible reasoning, by which he deduced Bayes’s rule. To apply Bayes’s rule we have to start from some priors, and Jaynes advocates the use of the maximum entropy principle to set up priors. However, there are many situations in which *both* MaxEnt and Bayesian inference are applicable. Which method should we take then? And do they return the same result? In this paper I propose a different set of axioms of plausible reasoning, by which I derive the minimum relative entropy principle¹, Bayes’s rule, and maximum likelihood.

^{*}Department of Economics, Yale University. Email: alexisakira.toda@yale.edu

[†]This paper benefited from conversations with Sylvain Barde and Sander Heinsalu. I am deeply indebted to my thesis advisor, Donald Brown. I thank Duncan K. Foley for introducing me to the works of Jaynes. The financial supports from the Cowles Foundation, the Nakajima Foundation, and Yale University are greatly acknowledged.

¹To the best of my knowledge, the minimum relative entropy principle was first introduced by Kullback [15, p. 37] under the name the principle of *minimum discrimination information*.

I proceed in two steps. First, I list the desiderata of a measure of information gain when a decision maker updates the plausibility of a proposition upon receiving new information. From these desiderata I derive the functional form of information gain. Second, I impose the decision maker to be maximally conservative, given all the relevant information. That is, the decision maker updates the plausibilities by minimizing the average information gain (*i.e.*, sticks to his or her prior as much as possible) subject to all relevant information, which I call the *Optimum Information Principle*. I show that the Optimum Information Principle implies the well-known minimum relative entropy principle, the Bayes rule, and also Jaynes’s axioms.

2 Axioms of Plausible Reasoning

Viewing probability as the plausibility of a proposition dates back at least to Keynes [14]. As Cox [3] describes it “as if Euclid had placed the Pythagorean theorem among the axioms of plane geometry”, Keynes’s axioms were not fundamental, and have been improved by [13] and [3]. To date the most primitive axioms of plausible reasoning seem to be those of Jaynes [12, pp. 17–19]:

J-I. Degrees of plausibility are represented by real numbers.

J-II. Qualitative correspondence with common sense.

J-III. Consistency.

- (a) If a conclusion can be reasoned out in more than one way, then every possible way must lead to the same result.
- (b) The robot² always takes into account all of the evidence it has relevant to a question. It does not arbitrarily ignore some of the information, basing its conclusions only on what remains. In other words, the robot is completely nonideological.
- (c) The robot always represents equivalent states of knowledge by equivalent plausibility assignments. That is, if in two problems the robot’s state of knowledge is the same (except perhaps for the labeling of the propositions), then it must assign the same plausibilities in both.

Desideratum II means the following. If we denote the plausibility of a proposition A given information I by $p(A|I)$, then

$$p(A|C') > p(A|C) \implies p(\neg A|C') < p(\neg A|C), \text{ and} \quad (2.1)$$

$$\left. \begin{array}{l} p(A|C') > p(A|C) \\ p(B|A \wedge C') = p(B|A \wedge C) \end{array} \right\} \implies p(A \wedge B|C') \geq p(A \wedge B|C). \quad (2.2)$$

In words, (2.1) says that if information C gets updated to C' in such a way that the plausibility of A is increased, then the plausibility of the negation of A is decreased; (2.2) says that if, in addition, the plausibility of B given A is unchanged, then the plausibility that both A and B are true must increase.

²The “robot” is a machine that performs plausible reasoning according to the desiderata.

Chapter 2 of [12] shows that desiderata I–IIIb imply that plausibilities have a probability representation and they obey Bayes’s rule, and that desideratum IIIc implies Laplace’s Principle of Indifference [17] for setting up priors.

In order to derive MaxEnt and Bayes’s rule, I first axiomatize the quantity which I call *information gain* and derive its functional form. The axioms, which are all intuitively appealing, are as follows.

- IG-1. Numerical representation: the information gain I is a function of prior plausibility p and posterior plausibility q .
- IG-2. Smoothness and monotonicity: the information gain is a smooth, increasing function in posterior plausibility.
- IG-3. Path independence: the total information gain of updating the prior plausibility p to the posterior q is independent of the path it is updated. That is, if there are two paths $p \rightarrow r \rightarrow q$ and $p \rightarrow r' \rightarrow q$, then $I(p, r) + I(r, q) = I(p, r') + I(r', q)$.
- IG-4. Independence from the choice of unit: whatever unit we choose to describe plausibility, the information gain should have the same value. That is, $I(tp, tq) = I(p, q)$ for $t > 0$.
- IG-5. Zero information gain for not updating: for any p , we have $I(p, p) = 0$.

Proposition 1. *Suppose that axioms IG-1–IG-5 hold. Then $I(p, q) = k \log \frac{q}{p}$, where $k > 0$ is an arbitrary constant.*

Proof. Since by axiom IG-2 the information gain $I(p, q)$ is smooth in q , it is partially differentiable with respect to q and I can be recovered by integrating its partial derivative. Differentiating $I(p, r) + I(r, q) = I(p, r') + I(r', q)$ with respect to q , we get

$$\frac{\partial I}{\partial q}(r, q) = \frac{\partial I}{\partial q}(r', q). \quad (2.3)$$

The left-hand side of (2.3) is a function of (r, q) , and the right-hand side of (2.3) is a function of (r', q) . Since r, r' are arbitrary, (2.3) must be a function of only q . Let $\frac{\partial I}{\partial q}(r, q) = g(q)$. By integration we get $I(r, q) = F(r) + G(q)$, where F is some function and $G = \int g$. By the path independence axiom IG-3, we get

$$\begin{aligned} [F(p) + G(r)] + [F(r) + G(q)] &= [F(p) + G(r')] + [F(r') + G(q)] \\ \iff F(r) + G(r) &= F(r') + G(r'). \end{aligned} \quad (2.4)$$

Since (2.4) holds for any r, r' , $F(r) + G(r)$ is constant, but it must be zero by axiom IG-5: $F(r) + G(r) = I(r, r) = 0$. Therefore $I(p, q) = F(p) + G(q) = G(q) - G(p)$. By axiom IG-4, we have $G(tq) - G(tp) = G(q) - G(p)$. Differentiating both sides with respect to q , we get $tG'(tq) = G'(q)$. Multiplying both sides by q and letting $x = tq$, we get $xG'(x) = qG'(q)$, so the function $xG'(x)$ is a constant k . Integrating $G'(x) = k/x$ yields $G(x) = k \log x + C$, hence $I(p, q) = G(q) - G(p) = k \log \frac{q}{p}$. Since I is increasing in q by axiom IG-2, we get $k > 0$. Clearly this function satisfies all axioms IG-1–IG-5. \square

From now on let us normalize the arbitrary constant k to 1, so the information gain is given by $I(p, q) = \log \frac{q}{p}$. This result,

$$\text{information gain} = \log \frac{\text{posterior plausibility}}{\text{prior plausibility}},$$

is mathematically identical to [5, p. 4], although Goldman takes this as the definition.³

In order to make plausible reasoning based on available information, consider the following desiderata.

- I. Degrees of plausibility are represented by probabilities.
- II. The robot always takes into account all of the evidence it has relevant to a question. It does not arbitrarily ignore some of the information, basing its conclusions only on what remains. In other words, the robot is completely nonideological.
- III. Aristotelian logic: the robot assigns zero plausibility to propositions that contradict its knowledge.
- IV. The robot always represents equivalent states of knowledge by equivalent plausibility assignments. That is, if in two problems the robot's state of knowledge is the same (except perhaps for the labeling of the propositions), then it must assign the same plausibilities in both.
- V. Given prior plausibilities, the robot updates the plausibilities by minimizing the average information gain of the posterior plausibilities subject to known information. In other words, the robot is maximally conservative.

Desideratum I is stronger than Jaynes's desideratum J-I because I assume that the plausibility is a probability (*i.e.*, finitely or countably additive measure). In particular, the plausibilities of mutually exclusive propositions are additive: if A, B are mutually exclusive propositions, then $p(A \vee B) = p(A) + p(B)$. Desideratum II is identical to J-IIIb. Desideratum III might be interpreted as a special case of II and probably needs no justification, but I need it nevertheless. Desideratum IV is identical to J-IIIc, Laplace's Principle of Indifference, which may or may not be necessary to prove subsequent theorems.

Desideratum V is the major difference from Jaynes's axioms. While Jaynes imposes "qualitative correspondence with common sense" (J-II), I impose that the robot is maximally conservative. This axiom makes sense, for if the robot radically updates the plausibilities (*i.e.*, not sticking to its prior), then it should not have set up the particular prior plausibilities in the first place. To avoid unnecessary reference to axiom numbers, let us group the desiderata as follows:

³In information theory the quantity $-\log p$ is known as the *self-information*, although I was unable to find a reference for its origin (Tribus [21] calls it *surprisal*). Our information gain $I(p, q) = \log \frac{q}{p}$ is the difference of the self-information of the prior and posterior. Kullback and Leibler [16] call $\log \frac{p_1}{p_2}$ the *information for discrimination*, where p_1, p_2 are general probabilities and not necessarily the prior and the posterior. The prior/posterior interpretation of p and q can also be clearly seen in [7, 8].

I–III:	Weak Axioms of Plausible Reasoning
I–IV:	Strong Axioms of Plausible Reasoning
IG-1–IG-5 and V:	Minimum Information Gain Principle

3 Implications of the Axioms

In this section I show that the new axioms imply Bayesian inference, maximum likelihood, maximum entropy principle, and minimum relative entropy principle.

Theorem 2. *Weak plausibility and the minimum information gain principle imply the minimum relative entropy principle (the minimum discrimination information principle of Kullback [15, p. 37]).*

Proof. Let $\{A_i\}$ be propositions that are mutually exclusive and exhaustive. Let $p_i = p(A_i|I)$ be the prior plausibility of proposition A_i given background information I , and $q_i = p(A_i|I')$ be the posterior plausibility to be computed given the new information I' . By desideratum I, we have $p_i, q_i \geq 0$ and $\sum p_i = \sum q_i = 1$. Since by Proposition 1 the information gain of A_i is $\log \frac{q_i}{p_i}$, the *ex post* average information gain is

$$H(q;p) := \sum_{i=1}^n q_i \log \frac{q_i}{p_i},$$

the relative entropy.⁴ By desiderata II and V, the robot minimizes $H(q;p)$ subject to all known information I' and the constraints $q_i \geq 0$, $\sum q_i = 1$, which is precisely the minimum relative entropy principle. \square

Corollary 3. *Strong plausibility and the minimum information gain principle imply the maximum entropy principle of Jaynes [9] for setting up priors.*

Proof. Desideratum IV is nothing but Laplace’s Principle of Indifference. Hence, by desideratum I, the robot assigns the prior plausibility $p(A_i) = \frac{1}{n}$. By Theorem 2 the robot computes the posterior plausibility $p_i = p(A_i|I)$ by minimizing

$$\sum_{i=1}^n p_i \log \frac{p_i}{1/n} = \sum_{i=1}^n p_i (\log p_i + \log n) = \sum_{i=1}^n p_i \log p_i + \log n,$$

(where we have invoked desideratum I: $\sum p_i = 1$) or equivalently, maximizing Shannon’s entropy $H(p) = -\sum_{i=1}^n p_i \log p_i$ [19]. This is precisely Jaynes’s maximum entropy principle [9]. \square

I propose to define the Optimum Information Principle by the combination of the weak or strong plausibility and the minimum information gain principle, despite its implication is the well-known minimum relative entropy principle. There are two reasons to avoid the term “entropy”. First, “entropy” is a misnomer both in physics (see [2]) and in information theory. According to [22], Shannon [19] named his measure of uncertainty or missing information “entropy” following the advice of von Neumann: “[It] has been used in statistical mechanics under that name . . . [and] no one knows what entropy really is, so in

⁴This quantity was first proposed by Kullback and Leibler [16], which they call, appropriately, “information”.

a debate you will always have the advantage.” Clausius coined the word “entropy” after the Greek word for “transformation”; given that “entropy” is a misnomer, adding the adjective “relative” makes it only worse. Second, as a measure of information gain the Kullback-Leibler information $H(q;p)$ is more fundamental than the Shannon entropy $H(p)$ as shown by the above axiomatic derivation as well as the comparison of the two information measures provided in [8]: the Kullback-Leibler information, unlike the Shannon entropy, extends to arbitrary probability measures and it satisfies an additivity property. Since by desideratum V the quantity $H(q;p)$ is the average information gain, and since Kullback and Leibler [16] call $H(q;p)$ “information” before the term “relative entropy” was coined, the term Optimum Information Principle seems best.⁵

Theorem 4. *Weak plausibility and the minimum information gain principle imply Jaynes’s desiderata I–IIIb, in particular Bayes’s rule. Therefore, the Optimum Information Principle is consistent with Bayesian inference.*

Proof. Let us first prove Bayes’s rule. Suppose that the robot is given background information I and that the robot has prior plausibilities on the propositions A_1, \dots, A_n, B , and any logical conjunction or negation generated by them. Therefore the prior plausibilities of $A_i \wedge A_j$, $A_i \wedge B$, $A_i \wedge (\neg B)$, etc., which are denoted by $p(A_i \cap A_j|I)$, $p(A_i \cap B|I)$, $p(A_i \cap B^c|I)$, etc., are well defined. The task of the robot is to update the plausibilities of $\{A_i\}$ when it is given additional information B . Since there are only a finite number of propositions, without loss of generality we may assume that $\{A_i\}$ are mutually exclusive and exhaustive. By desideratum I, we have $\sum_{i=1}^n p(A_i|I) = 1$.

Let us denote the posterior plausibilities by $q(A_i \cap B|B \cap I)$, etc. In order to compute them, by Theorem 2 the robot solves

$$\min_q \sum q \log \frac{q}{p} \quad \text{subject to} \quad (3.1a)$$

$$\sum_{i=1}^n (q(A_i \cap B|B \cap I) + q(A_i \cap B^c|B \cap I)) = 1, \quad (3.1b)$$

$$(\forall i) \quad q(A_i \cap B^c|B \cap I) = 0, \quad (3.1c)$$

where p, q in (3.1a) take all possible forms of $p(A_i \cap B|I)$, $q(A_i \cap B|B \cap I)$ and $p(A_i \cap B^c|I)$, $q(A_i \cap B^c|B \cap I)$. Conditions (3.1b) and (3.1c) come from desiderata I and III: since $\neg B$ (and hence $A_i \wedge (\neg B)$) is logically impossible knowing B , the robot assigns zero plausibility to $A_i \wedge (\neg B)$. That we impose (3.1b) and (3.1c) and nothing else comes from using all relevant information as in desideratum II.

Since the function $f(q) = q \log \frac{q}{p}$ is continuous and strictly convex and the constraints (3.1b), (3.1c) constitute a compact convex set, we can apply the Karush-Kuhn-Tucker theorem to solve (3.1). Let λ be the Lagrange multiplier corresponding to (3.1b). The Lagrangian is

$$L(q, \lambda) = \sum_{i=1}^n q_i \log \frac{q_i}{p_i} + \lambda \left(\sum_{i=1}^n q_i - 1 \right),$$

⁵[6] calls it the Maximum Information Principle, meaning that the missing information is maximized. Maximizing the missing information is equivalent to minimizing the information gain as we do here. The adjective “optimum” avoids the confusion between maximum/minimum.

where we have used the shorthand $q_i = q(A_i \cap B|B \cap I)$ and $p_i = p(A_i \cap B|I)$. The first-order condition, which is necessary and sufficient, reads

$$\frac{\partial L}{\partial q_i} = \log \frac{q_i}{p_i} + 1 + \lambda = 0.$$

This shows that q_i is proportional to p_i , so by $\sum q_i = 1$ we obtain $q_i = p_i / \sum_{i=1}^n p_i$. Therefore,

$$\begin{aligned} q(A_i|B \cap I) &= q(A_i \cap B|B \cap I) + q(A_i \cap B^c|B \cap I) \\ &= q(A_i \cap B|B \cap I) = q_i \\ &= \frac{p(A_i \cap B|I)}{\sum_{i=1}^n p(A_i \cap B|I)} = \frac{p(A_i \cap B|I)}{p(B|I)}, \end{aligned} \quad (3.2)$$

where the first equality holds because q is a probability (desideratum I), the second equality holds because $q(A_i \cap B^c|B \cap I) = 0$ (desideratum III), and the last equality holds because p is a probability and $\{A_i\}$ are mutually exclusive and exhaustive. (3.2) is precisely the Bayes rule.

Now let us show that Jaynes's desiderata I–IIIb are implied. All we need to show are desiderata II (conditions (2.1) and (2.2)) and IIIa. (2.1) holds because plausibility has a probability representation by desideratum I. (2.2) holds by Bayes's rule, which we have already deduced in (3.2). Desideratum IIIa holds by the Additivity Theorem of Hobson and Cheng [8, p. 308], where they essentially show that if the robot has initial background information I_0 that gets updated to I_1 and then to I_2 , with plausibilities p_0, p_1, p_2 respectively, then

$$H(p_2; p_0) = H(p_2; p_1) + H(p_1; p_0),$$

that is, the (minimized) Kullback-Leibler information is additive.⁶ In particular, if there are two ways to update, $I_0 \rightarrow I_1 \rightarrow I_2$ and $I_0 \rightarrow I'_1 \rightarrow I_2$, then we obtain

$$H(p_2; p_1) + H(p_1; p_0) = H(p_2; p'_1) + H(p'_1; p_0), \quad (3.3)$$

the path independence. Therefore Jaynes's desideratum IIIa holds because if a conclusion can be reasoned out in more than one way, the path independence property (3.3) ensures that every possible way leads to the same result. \square

At this point I stress the distinction between our axiomatization and other author's. In his seminal work [19], Shannon imposes as the third axiom “If a choice be broken down into two successive choices, the original H should be the weighted sum of the individual values of H ”, in which he implicitly uses Bayes's rule. The same remark applies to the axiomatization of the Kullback-Leibler information by Hobson [7]. Similarly, in the important axiomatization of the maximum entropy principle, Shore and Johnson [20] implicitly use Bayes's rule in their fourth axiom “Subset Independence: It should not matter whether one treats an independent subset of system states in terms of a separate conditional density or in terms of the full system density”. Zellner [24] derives Bayes's rule from an “information processing rule”, but it is not clear how it relates to maximum entropy and his definition of information seems somewhat arbitrary.

⁶This property is mathematically equivalent to the “Subset Independence” axiom of Shore and Johnson [20, p. 27].

On the contrary, Cox [3] and Jaynes [12] derive the Bayes rule from intuitively appealing first principles, as we have done. In addition we have derived the maximum entropy principle and the minimum relative entropy principle.

Finally, let us show that the Optimum Information Principle implies maximum likelihood.

Theorem 5. *The Optimum Information Principle implies the maximum likelihood principle of Fisher [4].⁷*

Proof. Suppose that $\{X_n\}_{n=1}^N$ are independently and identically distributed random variables with an unknown density f . Given the realizations $\{x_n\}$, suppose that the statistician wishes to fit a parametric density $f(x; \theta)$ to f , where $\theta \in \Theta$ is a parameter. Although prior and posterior distributions are meaningless for a frequentist, it is natural to interpret that the model $f(x; \theta)$ and the truth f correspond to the prior and posterior, respectively. Hence to make an optimal inference the statistician should choose θ so as to minimize the Kullback-Leibler information

$$H(f; f_\theta) = \int f(x) \log \frac{f(x)}{f(x; \theta)} dx.$$

However, by the law of large numbers we obtain

$$\begin{aligned} H(f; f_\theta) &= \int f \log f - \int f \log f_\theta = \int f \log f - E_f[\log f(X; \theta)] \\ &\approx \int f \log f - \frac{1}{N} \sum_{n=1}^N \log f(x_n; \theta), \end{aligned}$$

so the statistician should maximize the log likelihood $\sum_{n=1}^N \log f(x_n; \theta)$. \square

4 Concluding Remarks

The maximum entropy principle has occasionally been criticized *ad hoc* as “Why maximize entropy (or minimize relative entropy), why not other functions?”. An inference method is valuable if and only if it is useful in analyzing real data, and therefore an inference method requires no interpretation, and no justification except practical usefulness. (Nevertheless the justification of the maximum entropy principle has been provided [9, 10, 11, 20].) It is well-known that the minimum relative entropy principle (maximum entropy principle) and Bayesian inference are useful (see [18, 12] and the references therein). Therefore, since our Optimum Information Principle implies the minimum relative entropy principle, Bayes’s rule, as well as maximum likelihood, it should be equally useful. In addition we have axiomatized plausible reasoning and derived the maximum entropy principle; hence we have answered the question “why maximize entropy?”

References

- [1] Rev. Thomas Bayes. An essay toward solving a problem in the doctrine of chances. *Royal Society Philosophical Transactions*, 53:370–418, 1763.

⁷Although maximum likelihood is attributed to Fisher, it was already used by Laplace and Gauss a century before.

- [2] Arie Ben-Naim. *A Farewell to Entropy: Statistical Thermodynamics Based on Information*. World Scientific, Hackensack, NJ, 2008.
- [3] Richard T. Cox. Probability, frequency, and reasonable expectation. *American Journal of Physics*, 14(1):1–13, 1946.
- [4] Ronald A. Fisher. On an absolute criterion for fitting frequency curves. *Messenger of Mathematics*, 41:155–160, 1912.
- [5] Stanford Goldman. *Information Theory*. Prentice & Hall, Upper Saddle River, NJ, 1953.
- [6] Hermann Haken. *Synergetics: Introduction and Advanced Topics*. Springer, Berlin, 2004.
- [7] Arthur Hobson. A new theorem of information theory. *Journal of Statistical Physics*, 1(3):383–391, 1969.
- [8] Arthur Hobson and Bin-Kang Cheng. A comparison of the Shannon and Kullback information measures. *Journal of Statistical Physics*, 7(4):301–310, 1973.
- [9] Edwin T. Jaynes. Information theory and statistical mechanics, I. *Physical Review*, 106(4):620–630, 1957.
- [10] Edwin T. Jaynes. Prior probabilities. *IEEE Transactions on Systems Science and Cybernetics*, SSC-4(3):227–241, 1968.
- [11] Edwin T. Jaynes. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*, 70(9):939–952, 1982.
- [12] Edwin T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, U.K., 2003. Edited by G. Larry Bretthorst.
- [13] Harold Jeffreys. *Theory of Probability*. Clarendon Press, Oxford, 1939.
- [14] John Maynard Keynes. *A Treatise on Probability*. Macmillan, London, 1921.
- [15] Solomon Kullback. *Information Theory and Statistics*. John Wiley & Sons, New York, 1959.
- [16] Solomon Kullback and Richard A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [17] Pierre Simon Laplace. *Théorie Analytique des Probabilités*. Courcier, Paris, 1812.
- [18] Thomas J. Loredo. From Laplace to supernova SN 1987A: Bayesian inference in astrophysics. In P. F. Fougère, editor, *Maximum Entropy and Bayesian Methods*, pages 81–142. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1990.
- [19] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.

- [20] John E. Shore and Rodney W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, IT-26(1):26–37, 1980.
- [21] Myron Tribus. *Thermostatistics and Thermodynamics : An Introduction to Energy, Information and States of Matter, with Engineering Applications*. Van Nostrand, New York, 1961.
- [22] Myron Tribus and Edward C. McIrvine. Energy and information. *Scientific American*, 224(3):179–188, 1971.
- [23] Jan M. Van Campenhout and Thomas M. Cover. Maximum entropy and conditional probability. *IEEE Transactions on Information Theory*, IT-27(4):483–489, 1981.
- [24] Arnold Zellner. Optimal information processing and Bayes’s theorem. *American Statistician*, 42(4):278–280, 1988.